CENTRE DE RECHERCHE
SUR LES TRANSPORTS
(C.R.T.)
CENTRE FOR RESEARCH
ON TRANSPORTATION

UNIVERSITÉ DE MONTRÉAL
ÉCOLE DES HAUTES ÉTUDES COMMERCIALES
ÉCOLE POLYTECHNIQUE

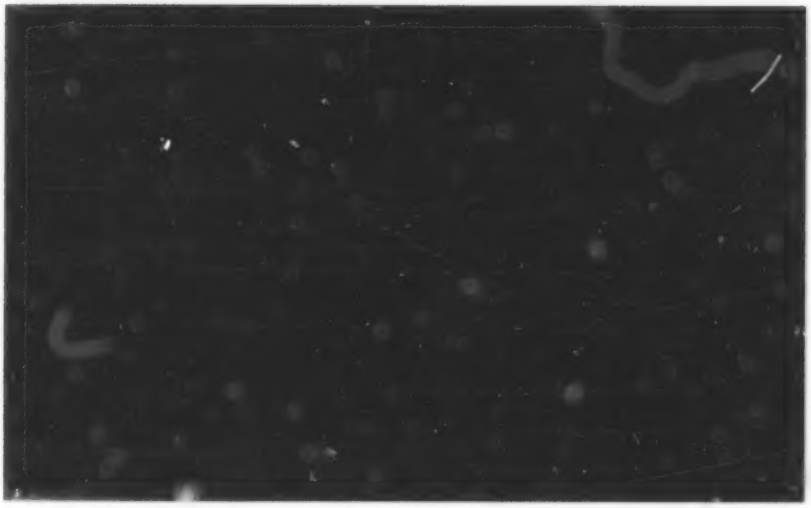

CENTRE DE RECHERCHE SUR LES TRANSPORTS
UNIVERSITÉ DE MONTRÉAL
C.P. 6128, SUCCURSALE A,
MONTRÉAL
CANADA H3C 3J7

TÉLÉPHONE: (514) 343-7575
TÉLEX: 05-24146 BIBPOLYTEC
EMAIL: crt@crt.umontreal.ca
TÉLÉCOPIEUR (FAX): (514) 343-7121

90-06645/

CRT-689

DESIGN OF REGULAR INTERCITY DRIVER ROUTES FOR THE LTL MOTOR-CARRIER INDUSTRY

by

Teodor Gabriel Crainic[1,2]
Jacques Roy[2]

[1] Centre de recherche sur les transports, Université de Montréal, C.P. 6128, Succursale A, Montréal (Québec) Canada H3C 3J7.

[2] Département des Sciences Administratives, Université du Québec à Montréal, C.P. 6192, Succursale A, Montréal (Québec) Canada H3C 4R2.

Centre de recherche sur les transports - Publication #689

May 1990

Dépôt légal - 2e trimestre 1990

ABSTRACT

The load plan of a LTL motor carrier determines the services it operates and the frequencies it offers. To satisfy the stable part of the demand for transportation, regular driver routes have to be built. These routes must cover the planned services and abide by the laws, regulations and conventions that regulate the drivers' workload, while guaranteeing the timely delivery of goods at the destination terminal. The objective is to minimize costs and to insure an efficient fleet utilization. We examine the problem characteristics and propose a modeling and algorithmic approach, based on set covering ideas. Results of an actual application are also presented.

KEY WORDS: LTL Motor-Carrier Transportation,
Routing and Crew Scheduling,
Regular Driver Routes,
Set Covering.

RESUME

Le plan de transport d'une entreprise de camionnage de charges partielles détermine les services et fréquences à offrir. Pour assurer le service de base et répondre à la partie stable de la demande, il faut établir un ensemble de routes–horaires régulières qui couvrent les services et qui respectent les lois, règlements et conventions régissant le travail des chauffeurs, tout en assurant la livraison à temps des marchandises au terminus de destination. L'objectif est de minimiser les coûts d'opération et d'assurer une utilisation efficace de l'équipement. Nous examinons les caractéristiques du problème et présentons une approche de modélisation et de résolution basée sur des techniques de recouvrement d'ensembles. Les résultats des expérimentations effectuées à partir d'une application industrielle sont également présentés.

MOTS CLES: Transport routier de marchandises en charges partielles,
Etablissement de routes et horaires,
Routes–horaires régulières,
Modèles et méthodes de recouvrement.

I. INTRODUCTION

Less-than-truckload (LTL) motor carriers often design *regular driver routes* in order to satisfy this part of the intercity freight transportation demand which is known to have a high degree of certainty. This relatively stable part of the demand for LTL services may account for approximately 60% to 80% of the total demand, according to seasonal variations. Then, carriers are able to provide their customers with a reliable service, while offering their drivers a choice among a set of regular routes. In many cases, according to their seniority in the company, drivers actually bid for the routes they will be driving for the next six months or so. The use of such regular routes also enables motor carriers to improve their control over the utilization of the equipment and, normally, facilitates the dispatching operations.

In order to properly design such regular routes, it is usually necessary to prepare ahead of time a transportation (load) plan that will take into account the ever present trade-offs between customer service performance and operating costs at a tactical planning level. This *transportation plan* generally determines:

- the type of transportation services to be operated between each pair of terminals and their frequencies;
- the itineraries (sequence of services and terminals) that are to be used to move freight from its origin to its destination;
- the workload, in terms of tonnage handled and number of trucks loaded and unloaded, at each terminal in the network;
- the necessary empty vehicle movements required to establish network balance.

Regular driver routes are then designed in such a manner as to cover those services derived from the tactical transportation plan, while respecting both government and labour contract rules. In this case, the objective is the on-time delivery of freight to its destination, while minimizing empty movements and operating costs.

The major difficulty nevertheless lies with the efficiency of the regular routes. Indeed, costs related to different routes will vary according to:

- the extent to which productive driver hours are actually used, the total amount of permissible production time being restricted by government regulation and labour contract rules;

- the amount of time required to hook up and change trailers;
- the time spent for general maintenance, fuel replenishment, various checks, etc.;
- the use of overtime;
- off-base allowances, such as meals, lodging, etc.

When solving such a difficult problem manually, managers will typically consider only a few alternatives in their analysis. They will most often be satisfied with a few improvements to the solution that was used during the previous planning period. The whole exercise will rest with the manager's experience and knowledge of the problem.

However, such an approach will not always provide motor carriers with effective and efficient regular driver routes. This is particularly true in the present context of rapid market changes and more restrictive safety regulations that affect drivers' time on the road and rest periods. Indeed, such regulation has just recently been enforced in the Province of Québec, thereby severely restricting the permissible length of direct links when designing driver routes.

We propose a new decision-support system that will facilitate the design of both effective and more efficient regular driver routes for the LTL motor-carrier industry. The system integrates operations research based procedures that insure an efficient and "optimal" generation and selection of regular driver routes, while also displaying interactive capabilities that allow the manager to inject his experience and knowledge into the process.

The objective of this paper is to present the general design of this decision-support system, and the modeling and algorithmic approach that we propose. The emphasis is on the formulation of a model which may adequately reflect the operational and decisional structure of the LTL motor-carrier industry, and which may be solved efficiently. We also demonstrate that even a simple utilization of our approach yields substantial improvements in the performance of a motor carrier's operations.

The paper is organized as follows. Section II presents the design of the decision-support system, while in Section III we specify the notation and some necessary definitions and discuss various modeling issues along with their algorithmic implications. Section IV is dedicated to the problem of generating segments. Section V presents the modeling of route productivity measures. Section VI deals with route generating procedures. The model and algorithm used to select the best set of regular driver routes are presented in Section VII. Section

VIII presents results obtained with real data, and we conclude with a discussion of further research and development issues in this area.

II. DESIGN OF THE DECISION-SUPPORT SYSTEM

Figure 1 shows a schematic representation of the planning process which leads to the design of regular intercity driver routes in the LTL motor-carrier industry. This planning process is implemented, under various forms, by most motor carriers which use regular driver routes. In our approach, we take a system view of the process, proposing analytical and computer tools integrated into a comprehensive decision-support system. In this section we only sketch the main lines of this system, the paper being focussed on the model and solution method that we propose for the design of regular driver routes.

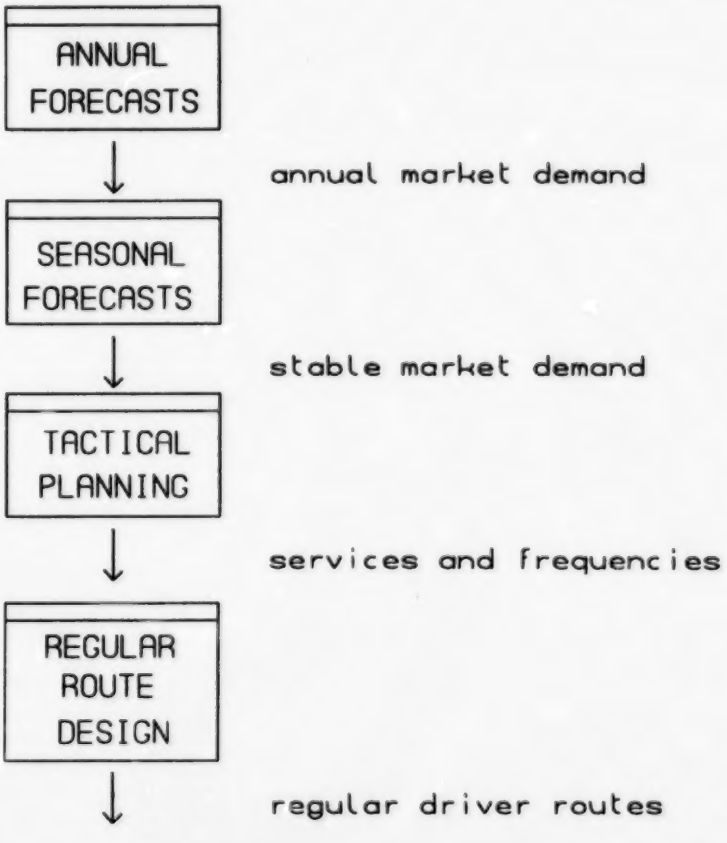

Figure 1. General Planning Approach for the Design of Regular Intercity Driver Routes

First, forecasts must be made for that part of the demand for transportation services which is considered to be stable over a given planning horizon (or season). Indeed, regular driver routes are only useful if a given level of demand can be determined with a high degree

of certainty. For major Canadian LTL carriers, this stable part of demand may represent from 60% to 80% of total demand, depending upon the season.

Forecasts must be made for each market served by the carrier. A *market* is determined by the origin and destination of the traffic, normally a pair of cities (terminals). Since we are only concerned with freight moving between existing terminals within a given motor carrier's network, we assume that demand is specified by the amount of freight to be moved from an origin to a destination terminal, for a given period of time.

These market forecasts are then used as inputs to the *tactical planning process* (Crainic and Roy, 1988), which produces a *transportation* (or *load*) *plan* for the stable part of demand for LTL services. Such a plan can be established by using the NETPLAN software (Roy and Delorme, 1989) which was developed from a tactical planning model proposed by Roy (1984) and based on a methodology put forward by Crainic (1982, 1988) and Crainic and Rousseau (1986).

Services and *itineraries* are important elements of the tactical planning process. A transportation *service* is determined by a transportation mode that can be used between two terminals within the network. A service is also defined by the *path* it follows through the physical network from its origin to its destination. Although this latter characteristic is often neglected at the tactical planning level, it plays a very important role when it comes to designing intercity driver routes, as we will see later on. The *frequency* is another important aspect of a service; it determines the number of departures, the *level of service*, associated with a particular transportation service for a given period. The set of all transportation services that a carrier makes available to its customers defines its *service network*.

For a given month, an *itinerary* defines how freight will flow from origin to destination. It determines the sequence of services used to move the freight, and the type of operation performed on freight at each intermediate terminal. Even though itineraries play an important role in defining the transportation plan, as do services, itineraries do not directly affect the design of regular driver routes; services, however, do.

Indeed, the type and frequencies of the transportation services used to satisfy the stable part of demand will be determined through the tactical planning process, and will then be used as input data to the regular driver route design process, as illustrated in Figures 1 and 2.

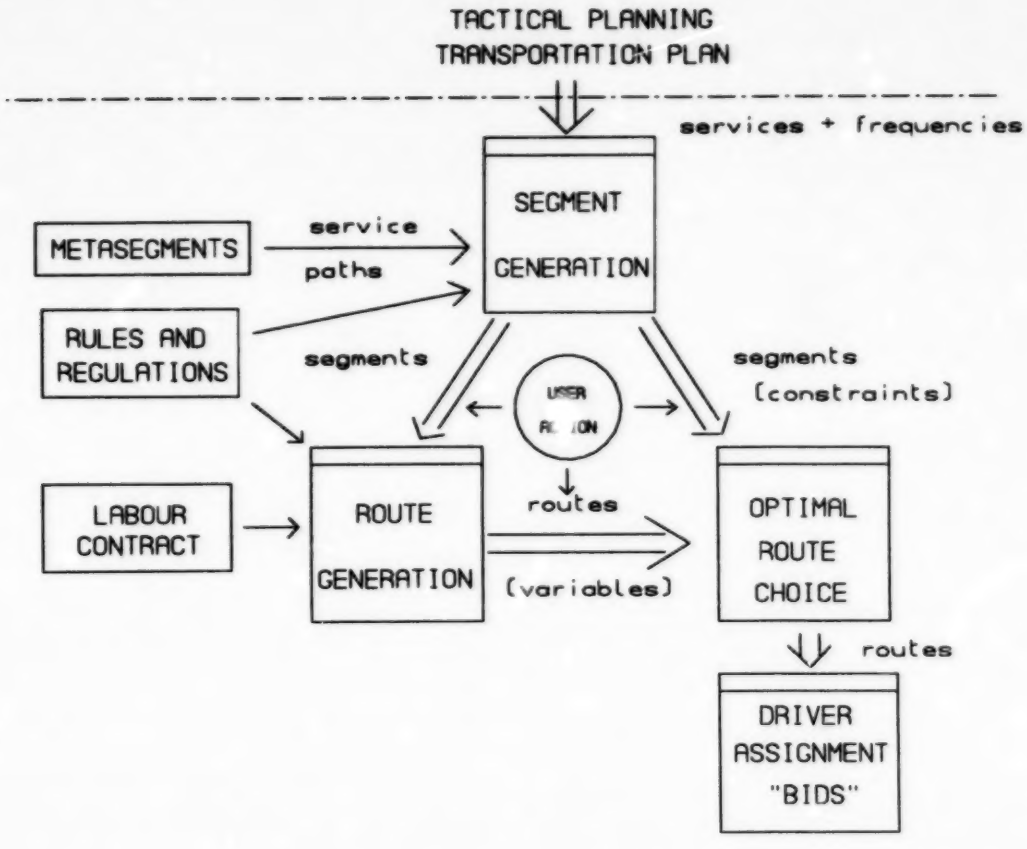

Figure 2. General Design of the Solution Approach

We have not found any published material dealing with the problem of the design of regular driver routes. An examination of the characteristics of the problem, especially in comparison with those of crew scheduling problems in other transportation modes (Rousseau, 1984), allows us, however, to identify certain similarities with the airline crew scheduling problem (Crainic and Rousseau, 1987). We have made use of these similarities in developing our solution method.

We propose a solution approach, illustrated in Figure 2, which starts with the transportation services and frequencies specified by the tactical load plan, and which takes into account both government and company rules and regulations. The problem is dealt with in three stages:

(i) *Segment Generation*. The services specified in the transportation plan must first be broken down into a set of *segments*. Each segment represents a direct driving link, and each service path yields one or several segments.

(ii) *Route Generation*. In order to satisfy the stable part of demand, the segments generated in the first stage of the procedure must be covered by regular driver routes. Such a set of efficient legal routes is generated at this stage.

(iii) *Route Selection*. We propose a model and solution approach which will select an optimal set of regular routes which satisfy the demand expressed by the segments generated above.

Routes from this optimal set are then assigned to drivers on a seniority basis as part of a bidding process. It is also important to note that we have accounted for the possibility of user interaction at each stage of the driver route design process. This would enable users to manage the overall process and, for example, bring various unquantifiable criteria into the process, or impose certain personal preferences which, otherwise, would go unobserved. User interaction will be made possible through the design of interactive-graphic planning systems which will include the models and solution methods presented in this paper.

III. DEFINITIONS AND NOTATIONS

Let us define the *planning period*, or simply the *period*, as the span of time, usually one week, for which regular driver routes have to be specified. These routes will then be repeated for all subsequent periods. Starting times for all components of the problem are computed with respect to the beginning of the planning period, which thus represents the "zero hour".

A period is usually broken down into several days. A *day* corresponds to a cycle of work and free time intervals (including rest). In intercity trucking, however, a working day does not normally coincide with a typical calendar day: a driver's workday starts in the evening and finishes the following morning. We should therefore refer to a driver's working period as a "work night". Yet, for clarity sake, we shall use the more conventional concepts of workday and rest periods.

Our approach to the problem is based on three major concepts: the *service path*, the *segment* and the *driver route*. These elements may be defined according to various attributes. Some of these are descriptive: location, length, etc. Others are time related and indicate when a particular movement occurs. We call this moment the *starting time*.

It is important to note that an element has only one physical description, but that it may have several time occurrences. For example, one could cover the same path several times during a given planning period. To distinguish between descriptive and time-related attributes, we use the prefix *meta* whenever the temporal attribute of an element is irrelevant.

Motor-carrier operations are best described in terms of various *networks*. These networks are all defined by a common set of nodes; the link sets may vary.

A *node* is a geographical location, where an operation may be performed. It is either a *terminal* or a *transfer point*. A trailer exchange may be performed at a transfer point, but no service may begin or end at such locations. A *base* is a terminal where drivers are stationed. Figure 3 provides an illustration of these definitions. Nodes 1, 2 and 3 represent terminals which are also bases, while node 4 is only a transfer point.

We use B as the set of all bases, T as the set of all terminals and N as the set of all nodes, with $B \subseteq T \subseteq N$. Define d_j^n as the *average delay* required to perform the specified operation at node j, $j \in N$. This delay will obviously depend on the type of the node.

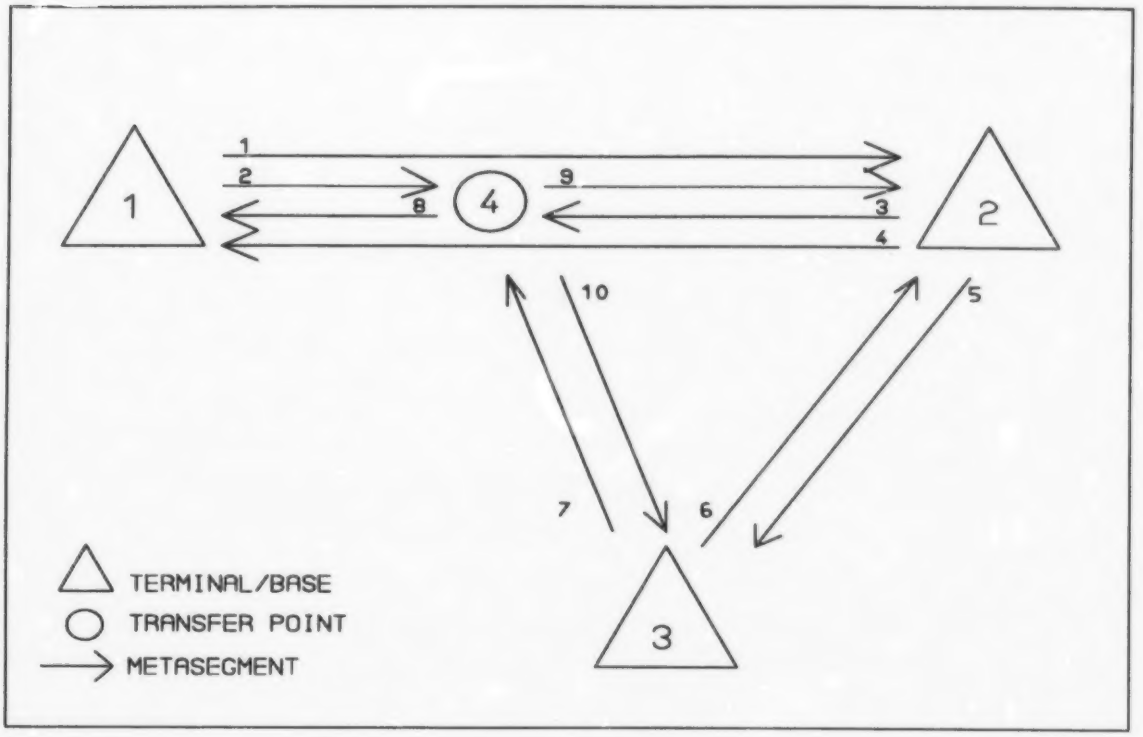

Figure 3. Network of Nodes and Metasegments

A *metasegment* represents a direct link between two nodes in the network which may be accomplished, according to company rules and regulations, without any intermediate operations. Let Σ be the set of all metasegments. For a given metasegment s, $s \in \Sigma$, we define:

- *Origin* O_s, $O_s \in N$;
- *Destination* D_s, $D_s \in N$;
- *Average driving time* d_s^p;
- *Frequency* f_s^p : the number of times that the metasegment is used by the services specified in the current transportation plan during the planning period.

The metasegments define a network, $G_\Sigma = (N, \Sigma)$, of permissible movements, according to company rules and regulations. For example, in Figure 3, we find that a total of ten metasegments determine the permissible movements between the nodes in the network. In this case, the absence of metasegments between nodes 1 and 3 does not necessarily reflect an absence of roads between those two points, but may possibly refer to a particular company policy.

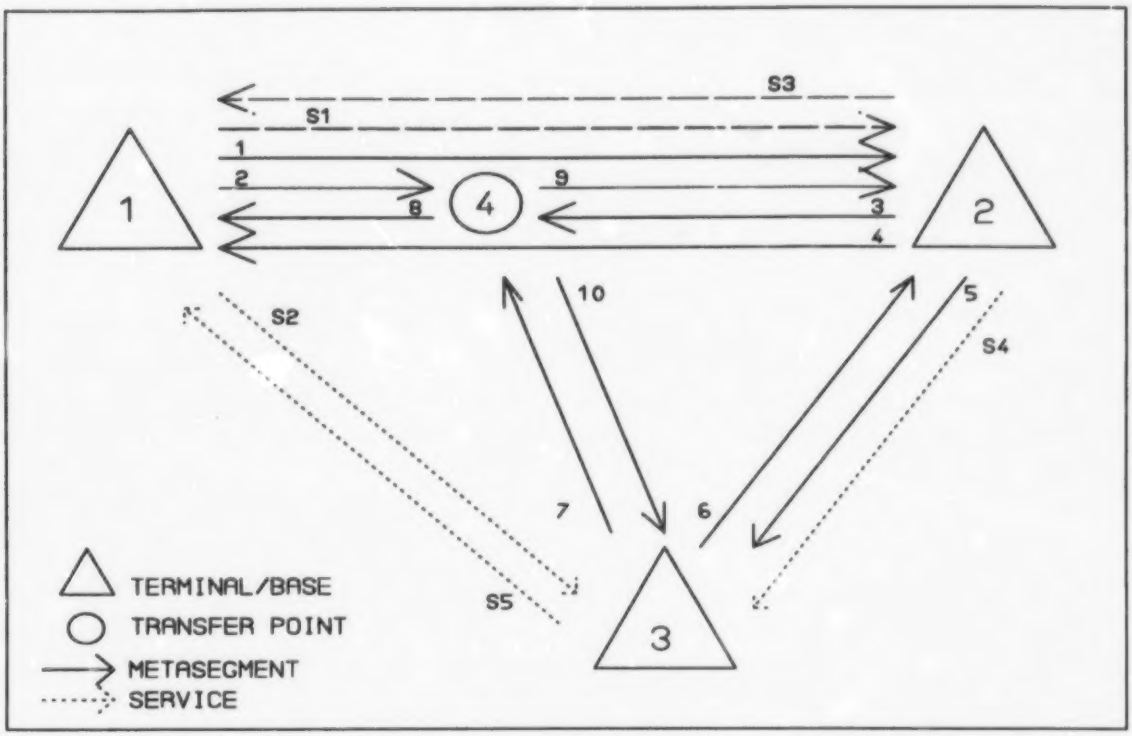

Figure 4. Network of Services and Metasegments

The metasegment network also determines the various service paths which may be used to implement the transportation plan. Thus, let us suppose that five services, illustrated in Figure 4, are to be operated on the network and covered by the metasegments that we have already shown. This would result in several service paths, as illustrated in Figure 5. In particular, we observe that a single service S1, for example, may use several paths, such as S1.1 or S1.2 and S1.0. On the other hand, we also note that a given metasegment may be part of several service paths, such as metasegment 2 for services S1 and S2, 3 for S3 and S4, etc.

A *segment* $s, s \in S$, is a metasegment to which a temporal dimension is added by the specification of a *starting time* t_s. Hence, several segments, such as (O_s, D_s, t_{s1}), (O_s, D_s, t_{s2}), ..., (O_s, D_s, t_{sn}), could correspond to a given metasegment (O_s, D_s) each segment individualized exclusively by its own starting time. S represents the set of all segments.

Segments and metasegments represent the basic units which make up the service paths and the driver routes. They therefore play a major role in our model (constraints are derived

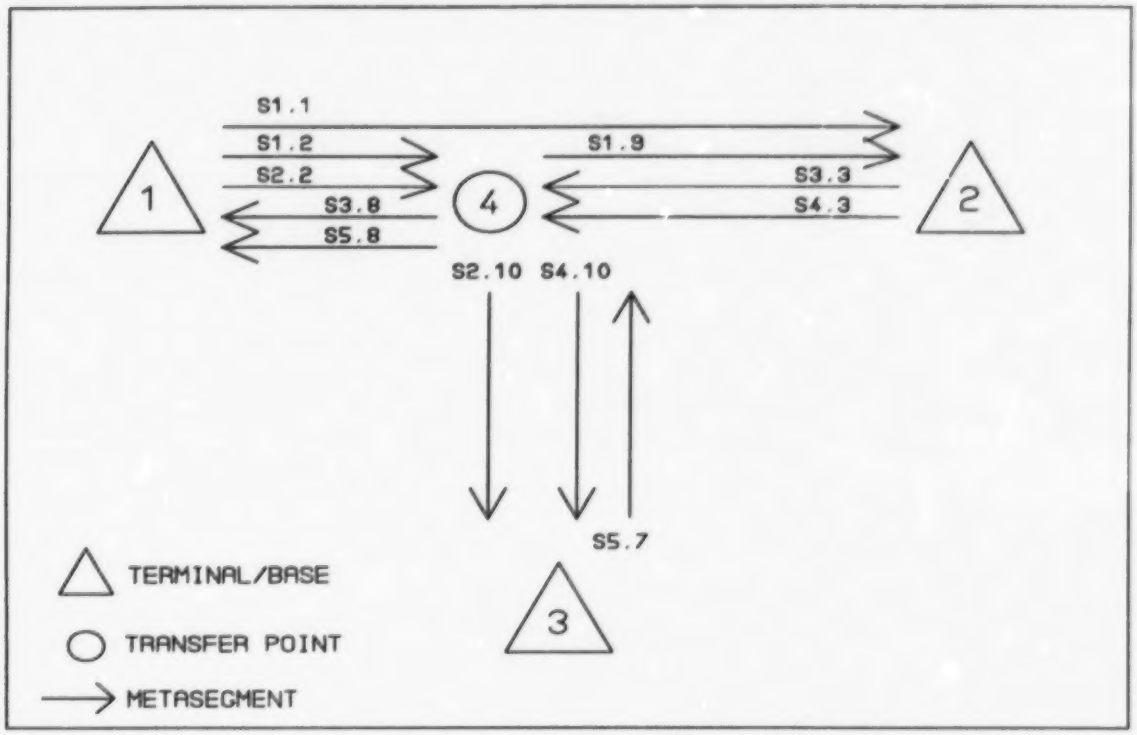

Figure 5. Network of Service Paths

from these elements) and in the procedures aimed at service analysis and generation of driver routes.

A *metaroute* is a legal sequence of metasegments, which starts and ends at the same base. The metaroute is said to be *normal* if it does not include any rest period outside its base, and *extended*, otherwise. A metaroute whose composing metasegments appear only once is called *simple*; otherwise, it is *cyclic*. Table 1 displays a few metaroutes built by using the metasegments illustrated in Figures 3 and 4. Note that metasegment (2,2) represents a rest period at node 2.

Several attributes define a metaroute r, $r \in \Omega$, where Ω is the set of all metaroutes. Those required by our formulation are:

- its *base* b_r, $b_r \in B \subseteq N$, which is both the origin and the destination of the metaroute;
- its *composition* κ_r, which is made up of a sequence of metasegments and, eventually, off-base rest periods;

DESCRIPTION	TYPE
(1,A) (4,3) (3,A) (4,1)	NORMAL – SIMPLE
(1,2) (2,2) (2,A) (4,1)	EXTENDED – SIMPLE
(1,A) (4,1) (1,A) (4,1)	NORMAL – CYCLIC

Table 1. Examples of Metaroutes

- the *total trip time* d_r^t, which includes the driving time, the waiting time at the nodes of the network, and the length of the rest periods;
- the *driving time* d_r^p, which is computed as the sum of the average driving times for all the metasegments that make up the metaroute:

$$d_r^p = \sum_{i \in \kappa_r} d_i^p ;$$

- the *duty time* d_r^w, which represents the amount of time the driver is available for work, e.g. while he is driving a vehicle or waiting for the completion of an operation at a node;
- its *cost* c_r, which is defined in Section V.

A *drive route* (or simply, a *route*) is a legal sequence of segments starting and ending at the same base. It is therefore a metaroute to which a time dimension has been added, by specifying the starting time of each constituent metasegment. Hence, a route has a specific *starting time*, t_r, which corresponds to the starting time of the first segment in its composition. With the appropriate adjustments, all other definitions and notations relative to metaroutes also apply to routes. In particular, the composition of a *cyclic* route displays at least two occurrences of the same metasegment, at different times.

These routes will then be assigned to drivers, according to their seniority in the company. Hence, they are the decision variables of our model.

IV. SEGMENT GENERATION

Segments are the basic elements of the whole process. Driver routes, metaroutes and service paths are all constructed by using segments, whether they are time related or not. Moreover, segments are also used to generate the constraints of the mathematical formulation proposed in Section VII to optimize the route selection.

Conceptually, segments are obtained from the service paths and frequencies specified in the (tactical) transportation plan. For every service departure, its path is broken down into metasegments, each of which generates a segment once its starting time is computed. This computation is done by adding up the driving time associated with each preceding metasegment and the delays at nodes to the time the service starts from its base.

In doing so, it is implicitly assumed that the starting time and exact path of each service departure are known ahead of time. However, this information is not always available. In fact, this data is not required for the preparation of the transportation plan at the tactical planning level. Furthermore, when it comes to implementing the transportation plan into day-to-day operations, dispatchers may significantly modify the service departure times, in order to respond to variations in the actual demand for transportation services. Therefore, when it comes to designing regular driver routes (a process which is half-way between tactical planning and day-to-day operations), one needs to come up with an estimate for the starting time, and an exact path, for each one of these service departures.

It is noteworthy that the absence of exact schedules for services, and hence segments, is one of the major differences between the problem of designing regular driver routes in the motor-carrier industry and the equivalent problem in the civil aviation industry.

Now, when a given service has only one feasible path in the metasegment network G_Σ, and only one departure planned during the period (frequency = 1), then the problem is reduced to a simple computation of segment departure times. The vast majority of services, however, have a frequency higher than one per planning period, besides having the possibility of following several paths in G_Σ. Paths are then selected according to various criteria; some are arbitrary, most are of an economical nature. In the case illustrated in Figure 4, for example, one notes that service S1 may be operated on two different paths: the first one uses metasegment 1 only, while the second travels on metasegments 2 and 9. The first alternative

is faster, but it forces the driver to spend a rest period off base. The second path is more convenient if we operate as follows: a driver leaves node 1 with a trailer for node 2, while another driver leaves node 2 with a trailer for node 1. At the assigned transfer point, here node 4, trailers are exchanged, and each driver heads for his own base, his tractor pulling the trailer from the other service. This way, the off-base rest periods are eliminated at the expense of a slight increase in the total transit time of each service and in the total duty time of each driver. In actual operations, the second alternative is used for early departures, thus permitting on time deliveries, while the second path has to be used for late departures if the target delivery times have to be met.

The segment generation problem may thus be thought of as the combined problem of *determining* and *decomposing* service paths, and of *distributing* service frequencies over time, while insuring on-time delivery of freight at destination, minimizing total times and reducing the amount of off-base rest periods.

In its general form, this problem is equivalent to the establishment of service schedules, that is, the simultaneous determination of service paths and starting times for all services. One possible approach to this problem requires that one works on a time-space dynamic network, in order to coordinate service arrivals and departures, and that the time-dependent distribution of each market demand forms the basic input to the process. Yet, we will see that this challenging but difficult problem need not be addressed exactly in order to design both effective and efficient driver routes.

Indeed, since we are working with the stable part of demand, the general pattern of freight arrival and service departures can be obtained from past experience. Furthermore, the exact coordination of daily departures is performed in real time by dispatchers, the driver being only advised ahead of time of a relatively wide time interval (in the order of three hours, according to Blanchette, 1989) within which his precise departure time will be determined.

We have therefore chosen to determine the starting times on the basis of data gathered from past experience. At the same time, we increase the flexibility of the route generator by taking into account the necessary coordination required to transfer trailers from one segment to another (Section VI). Our approach deals with each service separately and consists of two

steps. The first step determines the number of daily departures required according to the service frequency specified in the transportation plan for the given planning period. These departures are then distributed, for each day of the period, over a time interval which allows for freight to be handled prior to the first departure, and for the last departure to be early enough to insure on-time delivery at destination. In the first step, frequencies are usually distributed evenly over the number of working days of the planning period. For the second step, however, experience has shown that departures are not necessarily evenly spread over the interval; very early and very late departures may be observed, with most departures clustered during "peak" hours. In our experiment (see Section VIII), we tested a uniform distribution of departures, except for symmetrical metasegments (two metasegments linking the two same nodes, but in opposite directions), in which case the time interval was split in two, in order to account for possible connections between segments (Crainic, Roy and Delorme, 1989).

To determine which service paths, and hence metasegments, will be used, we propose a method which hinges on shortest path techniques. This approach is described in Crainic and Roy (1990), and allows an efficient access to a wider array of solutions than simply resorting to paths that were used historically by the company.

V. COST MEASURES FOR DRIVER ROUTES

"Cost" is the most widely used productivity measure for evaluating driver routes and the drivers' wages is the major and, often, the sole component. As we shall see, however, this measure is not sufficiently precise, nor sufficiently discriminant, for an optimization-based procedure which selects the best set of regular driver routes. We present in this section a productivity measure which, although based on the wages drivers receive, is better suited to the requirements of the route selection procedure.

Drivers' wages are normally regulated by labour contracts which guarantee, for each duty period, that a minimum number of duty hours shall be paid, irrespective of the actual amount of time worked. The wage formulae also specify the maximum amplitudes for normal duty and overtime, as well as the hourly rates for the two cases. A realistic example of such a formula might indicate that the equivalent of 9 worked hours are guaranteed for each duty period, that the maximum normal duty period is of 12 hours, including breaks, and that the driver may work a maximum of 14 hours, including overtime. Various other factors, such as distance rates (which vary according to the route), route dependent overcharges, and auxiliary expenses (e.g. monetary compensations for rest periods taken off-base), further complicate the accounting formula for driver costs. See Blanchette (1989) for a detailed presentation of the procedure used to compute drivers' wages for a Canadian LTL motor carrier. The similarity between this procedure and those used by North American airlines is remarkable.

A realistic representation of the *real* cost γ_r of a regular driver route is:

$$\gamma_r = \begin{cases} K^f & \text{if } 0 \le d_r^w \le t^f \\ K^f + c_r^n\,(d_r^w - t^f) & \text{if } t^f \le d_r^w \le t^n \\ K^f + c_r^n\,(t^n - t^f) + c_r^o\,(d_r^w - t^n) & \text{if } t^n \le d_r^w \le t^o \end{cases}$$

where

t^f : Minimum guaranteed duty time, per period;

t^n : Regular duty time amplitude for a period;

t^o : Maximum duty time in a period, including overtime;

K^f : Minimum guaranteed wage;

c_r^n : Hourly rate for normal duty time, for route r;

c_r^o : Hourly rate for overtime, for route r;

and the route is, of course, legal: $0 \leq d_r^w \leq t^o$.

Figure 6 illustrates this cost formula. Note that, for normal routes, the total trip time equals the total duty time, $d_r^t = d_r^w$. The cost of an extended route is computed as the sum of the costs of each of its duty periods, obtained by using this formula, plus the costs of the auxiliary expenses.

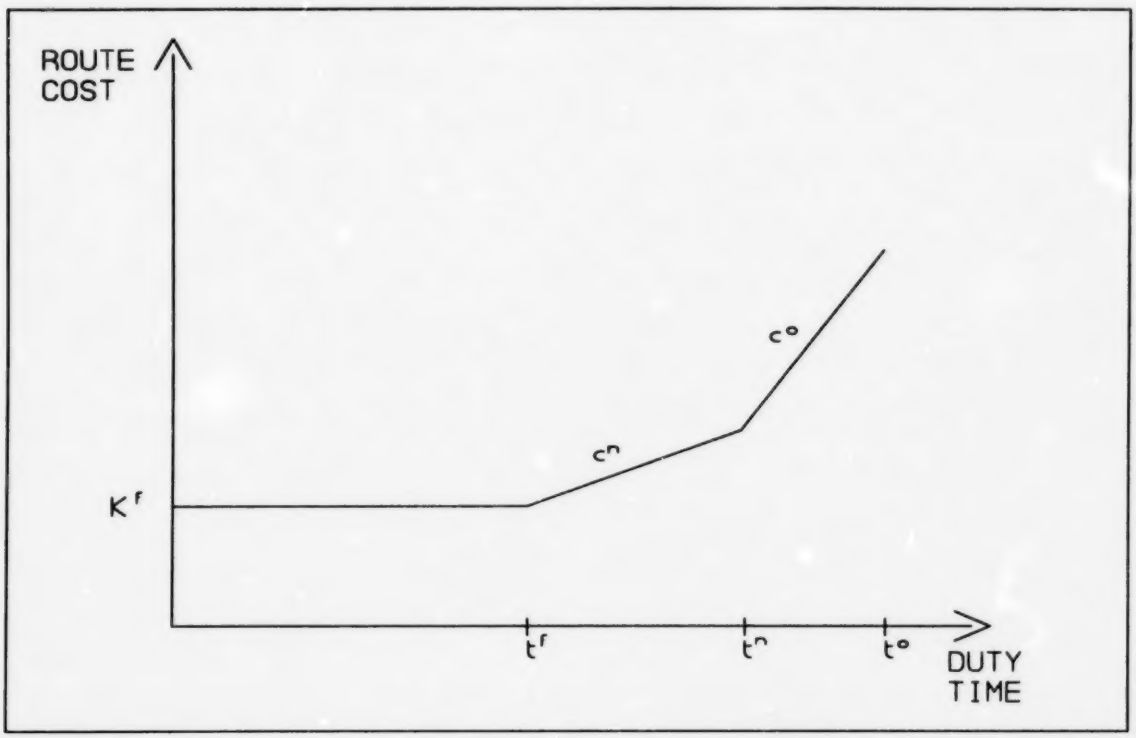

Figure 6. Cost Function for a Regular Driver Route

Two important factors, both related to the notion of *productivity* of a regular driver route, make this formula inappropriate for the computation of the objective function coefficients of an optimization-based route selection procedure. First, any route with a total trip time less than the guaranteed minimum receives the same cost, K^f, irrespective of the actual work done. Second, when using this type of formula, one cannot identify the operations that make up the cost of the route.

Yet, if one agrees that *driving* the truck represents the *productive* part of a driver's work assignment, one would prefer to select routes where a minimum number of duty hours correspond to non-driving activities. We examined several possible formulae for computing a route cost based on productivity measures. The results obtained by using each one of these alternatives are illustrated in Table 2 for six hypothetical routes. The total trip and driving times are indicated for each route, while the following constants have been used: $K^f = 90$, $t^f = 9$, $t^n = 12$, $t^o = 14$, $c_r^n = 10$, $c_r^o = 15$ (these constants are only indicative, they do not correspond to any actual application).

TRIP TIME	DRIVING TIME	ACTUAL COST	HYP 1	HYP 2	HYP 3a	HYP 3b	HYP 4
5hrs	4hrs	90.00	18.00	53.57	125.00	195.00	5.00
7hrs	4hrs	90.00	12.86	75.00	115.00	165.00	2.00
10hrs	7hrs	100.00	10.00	107.14	110.00	130.00	3.00
11hrs	7hrs	110.00	10.00	117.85	115.00	125.00	4.00
13hrs	10hrs	135.00	13.50	139.28	135.00	345.00	3.50
14hrs	10hrs	150.00	15.00	150.00	150.00	570.00	5.00

Table 2. Route Cost Measures – Various Formulae

Hypothesis 1.

Compute an *average hourly cost* of a route:

$$c_r = \frac{\gamma_r}{d_r^w}.$$

Hypothesis 2.

Compute an *average unit cost* over the maximum duty time length [0, t^o] :

$$c_r^m = \frac{K^f + c_r^n\,(t^n - t^f) + c_r^o\,(t^o - t^n)}{t^o}$$

which may afterward be used to compute the route cost $c_r = c_r^m\, d_r^w$.

Both preceding hypotheses and their variants are quite simple, even simplistic. These models require minimal computing efforts and generally solve the first problem (even if, as in the case illustrated in Table 2, they may not differentiate amongst routes with different actual costs). The productivity issue is, however, not addressed.

Hypothesis 3.

Specify an hourly rate K^i for the trip *idle time*, defined as the difference between the driving time and the normal amplitude of a duty period. The route cost is then:

$$\gamma_r = \begin{cases} K^f + K^i\ (t^n - d_r^w) & \text{if } 0 \le d_r^w \le t^f \\ K^f + c_r^n\ (d_r^w - t^f) + K^i\ (t^n - d_r^w) & \text{if } t^f \le d_r^w \le t^n \\ K^f + c_r^n\ (t^n - t^f) + c_r^o\ (d_r^w - t^n) & \text{if } t^n \le d_r^w \le t^o \end{cases} .$$

A penalty could also be applied to the overtime working hours.

We assume in this case that the normal duty period amplitude is a worthy goal for a productive driver route. This type of cost formula yields values which are not far from the actual wages paid to drivers and it also permits to penalize short and, when so desired, overtime routes. Moreover, the *idle time cost*, K^i, offers interesting opportunities for analysis concerning the relative value of the various activities that make up driver routes.

The calibration of the idle time cost represents, however, a very serious problem, since its value has a significant impact of the route costs and, therefore, on the final route choice. This phenomena is illustrated in Table 2, where the figures in column 3a correspond to a idle time value of $K^i = 5$, while those in column 3b were computed by using $K^i = 15$ and by penalizing the overtime hours by the same K^i factor. One may note the insidious effect of the idle time cost value on the relative cost ordering for routes with duty times between the minimum guaranteed and the normal amplitude. Moreover, this type of formulae does not adequately address the productivity issue, since two routes of equal total trip time will still receive the same cost, irrespective of their actual driving time.

Hypothesis 4.

Define the cost of a driver route as the time paid but non dedicated to the driving of the vehicle. The *paid time* for a regular route, d_r^s, may be defined as follows:

$$d_r^{\$} = \begin{cases} t^f & \text{if } 0 \le d_r^w \le t^f \\ d_r^w & \text{if } t^f \le d_r^w \le t^n \\ t^n + \dfrac{c_r^o}{c_r^n}\,(d_r^w - t^n) & \text{if } t^n \le d_r^w \le t^o \ . \end{cases}$$

Then, the cost of the route is $c_r = d_r^{\$} - d_r^p$.

For an n-duty-period extended route, the cost is computed by a similar formula. First, the paid time of each one of the n duty periods $d_i^{\$}$, $i = 1,...,n$, is computed by using the formula defined above. The cost of the extended route is then given by:

$$c_r = \sum_{i \in \kappa_r} d_i^{\$} + n\ K^{hb} - d_r^p.$$

The class of formulae described by this last hypothesis adequately addresses both the discrimination and the productivity issues. It assumes, of course, that time dedicated to activities other than driving the vehicle is not really productive for the carrier. Consequently, the most interesting set of regular driver routes is the one which minimizes this time.

To compute the route cost, all components and activities (e.g. overtime periods, off-base rest periods, etc.) that make up a driver route are first evaluated in terms of normal driving time. The conversion factor between normal and overtime working hours might vary; the ratio between the two corresponding cost rates seemed, however, appropriate, since it reflects the monetary relationships between the two activities. Then, the paid time is compared to the actual driving time, and the difference is assigned as the cost of the route. Thus, short routes are marked as "more expensive", as are those with extensive overtime periods and long waiting periods at terminals.

It is noteworthy that this formula displays characteristics which are similar to those of the formulae used to compute the crew pairing cost in most North American airlines (see Crainic and Rousseau, 1987; or Gershkoff, 1989). Similarly to those formulae, the procedures used to compute both the driver route costs and the associated productivity measures, as defined by the last two hypotheses, are intrinsically nonlinear: the cost variation due to the addition to, or the deletion of a segment from the path of a route, depends not only on the specific segment, but also on the position of the segment in the route path and on the route's

total trip and duty times. Consequently, one cannot evaluate the cost of a segment **prior** to the generation of the route, and the construction of normal routes is performed by using chaining procedures rather than network algorithms. These procedures are briefly discussed in the next section.

VI. GENERATING REGULAR DRIVER ROUTES

In the combinatorial sense, a driver route is a legal arrangement of segments, waiting periods at terminals and, eventually, off-base rest periods. The legal feasibility of a route is determined only by temporal criteria: total trip and duty times, time available to perform terminal activities, etc. Consequently, the total number of legal driver routes for a given case may be very high, especially when extended routes are contemplated. To generate all these routes would both be costly and algorithmically inefficient: the final set would be far too large for any efficient route selection procedure.

Our approach aims at generating only a limited number of "good" regular driver routes. This set is "good" in the sense that it contains the optimal solution of the route choice model presented in the next section. To achieve this objective, we decompose the generation procedure into several stages, that we execute, if necessary, at various points in the solution procedure.

The most fundamental decomposition occurs between the generation of normal and extended routes. Extended routes are, by far, the most numerous. Also, planners generally tend to avoid using this type of route since extended routes imply extra costs (e.g., for off-base rest periods). Additional arguments against the use of extended routes are: the relatively high density of the network of LTL terminals; the numerous, easy and inexpensive opportunities for implementing driver relief points (this is a major difference when compared to the airline crew scheduling problem); and, in Canada, the recent toughening up of the safety legislation concerning the duty and trip times. Consequently, only normal routes are usually generated. Extended routes are built only when required (e.g., when the regular services specified in the load plan cannot be covered by normal routes alone), in a post-optimal stage of the solution method (see Section VII).

Normal driver routes are generated by a *chaining* procedure, which chronologically examines the list of segments and combines them according to legal and efficiency criteria (e.g. maximum waiting time at a terminal). The procedure is described in detail by Crainic and Roy (1990). Here, we simply enumerate some of its most important features.

Normal routes are independently generated for each day of the planning period, and for each driver base in the network. Note that, in order to compensate for possible imprecisions

in the specification of the segment departure times, an *overlapping period* is allowed between the starting time of one segment and the arrival time of the previous segment in the route path. The temporal and economic characteristics of the route are, of course, correctly computed, the appropriate segment starting times being adjusted if the route belongs to the optimal solution.

To further limit the number of routes that are generated, the chaining process stops as soon as the initial base is returned to. Cyclic routes are then built by combining routes from this initial set. Feasibility rules, marginal values and post-optimality criteria are used to limit the number of initial routes considered in the process and that of new routes. See Crainic and Roy (1990) for details.

An extended driver route is composed of a sequence of segments, operations at terminals and off-base rest periods. Consequently, chaining procedures may be used to build them. Previous studies of various crew scheduling problems (Lavoie, Minoux and Odier, 1987; Crainic and Rousseau, 1987, etc.) have shown, however, that it is much more efficient to define the path of an extended route as a sequence of duty periods and off-base rest periods. Extended routes are then generated by a two-phase procedure:

(i) Generate duty periods;

(ii) Form extended routes by legally combining duty periods and rest periods.

The main advantage of this type of approach is that only a relatively small number of "good" routes are generated. Indeed, duty periods are still not very numerous (as compared to the total number of legal routes, of course), and are normally simultaneously generated with the normal routes. The route generation phase may then be postponed until quite late in the resolution process, and may be performed by a column generation type of algorithm. This approach has performed extremely well on a large gamut of crew scheduling problems (see, for example, Crainic and Rousseau, 1986; Crainic and Rousseau, 1987; Desrochers, 1986; Lavoie, Minoux and Odier, 1987; Rousseau, 1984; etc.).

Two methods, both developed for the airline crew scheduling problem, are available for the generation of extended driver routes. Lavoie, Minoux and Odier (1987) generate pairings by using a shortest path algorithm on an extended space-time diagram. The nodes of the diagram correspond to the various *states* that a duty period may be in, according to its

position in the route path and the current cost and productivity characteristics of the route, while links indicate a legal connection between two duty periods. Crainic and Rousseau (1987) take a different approach and use a chaining procedure, where strict rules (based on marginal values of pairings and duty periods) limit the access of duty periods to the generation process.

It is noteworthy that both methods require the same preliminary work (generation of duty periods), and take the same broad approach of solving a series of linear problems, each followed by computations of marginal costs and generation of routes. When using the first approach, however, one also has to enumerate the possible states of duty periods and build the extended diagram. Thus, this method performs well when the number of states is relatively low and the number of duty periods in a route is relatively high. On the other hand, the performance of the chaining approach becomes less interesting as the length, in number of duty periods, of the route path increases.

For LTL motor-carrier transportation, the number of possible states is relatively low, due to the limited number of rules (as compared with the air transport industry, for example) concerning the driver work load. Yet, the length of the typical route is also generally small, at least in the Canadian context: we have found only two-duty period routes in the data we have analyzed. We therefore use the chaining approach for the generation of extended regular driver routes.

VII. SELECTING THE OPTIMAL ROUTE SET

We propose a *set covering* approach in order to model the problem of selecting the best set of regular driver routes. The presentation of this model and of the associated algorithm is the subject of this section.

Consider R the set of all regular driver routes. With each normal route r, $r \in R$, we associate a cost c_r and a sequence of segments and, eventually, off-base rest periods, κ_r. We may rewrite this composition as a n_S-dimension vector, $a_r = (a_{ir}) \geq 0$, such as $a_{ir} = 1$ if segment i belongs to route r (and equals to 0 otherwise).

Let $x_r = 1$ if route r, $r \in R$, is chosen (and 0 otherwise). The model that we use to select the optimal set of regular driver routes may then be written as

$$
\begin{aligned}
(PRH) \qquad & \textit{Minimize} \quad \sum_{r \in R} c_r x_r \\
& \textit{Subjet to} \quad \sum_{r \in R} a_{ir} x_r \geq 1 \qquad \forall s \in S \\
& \qquad\qquad x_r \in \{ 0, 1 \} \qquad \forall r \in R \, .
\end{aligned}
$$

The model belongs to the now classical class of *set covering* formulations (Salkin and Mathur, 1989). Consequently, several reduction rules may be devised in order to decrease the size of the covering problem and thus increase the efficiency of the solution method. Crainic and Roy (1990) present several such rules and discuss their efficiency.

We propose a set covering formulation, rather than a set partitioning one, which insures that each segment is selected only once in the optimal set, since, in actual practice, one may observe movements of tractors without trailers. These movements are not very frequent, but they do occur, due to unbalances between supply and demand at various terminals of the network.

The main step on the way to determine the optimal solution to the above formulation consists in solving its linear relaxation:

$$(PRHL) \qquad \text{Minimize} \sum_{r \in R} c_r x_r$$

$$\text{Subjet to} \sum_{r \in R} a_{ir} x_r \geq 1 \qquad \forall s \in S$$

$$0 \leq x_r \leq 1 \qquad \forall r \in R.$$

To do this, we make intensive use of the very interesting structure displayed by the constraint matrix of the model (see Figure 7), where the columns correspond to routes and the rows to segments. Note that routes have been grouped in such a way that the normal routes for each day are shown together, and that segments have been sorted accordingly.

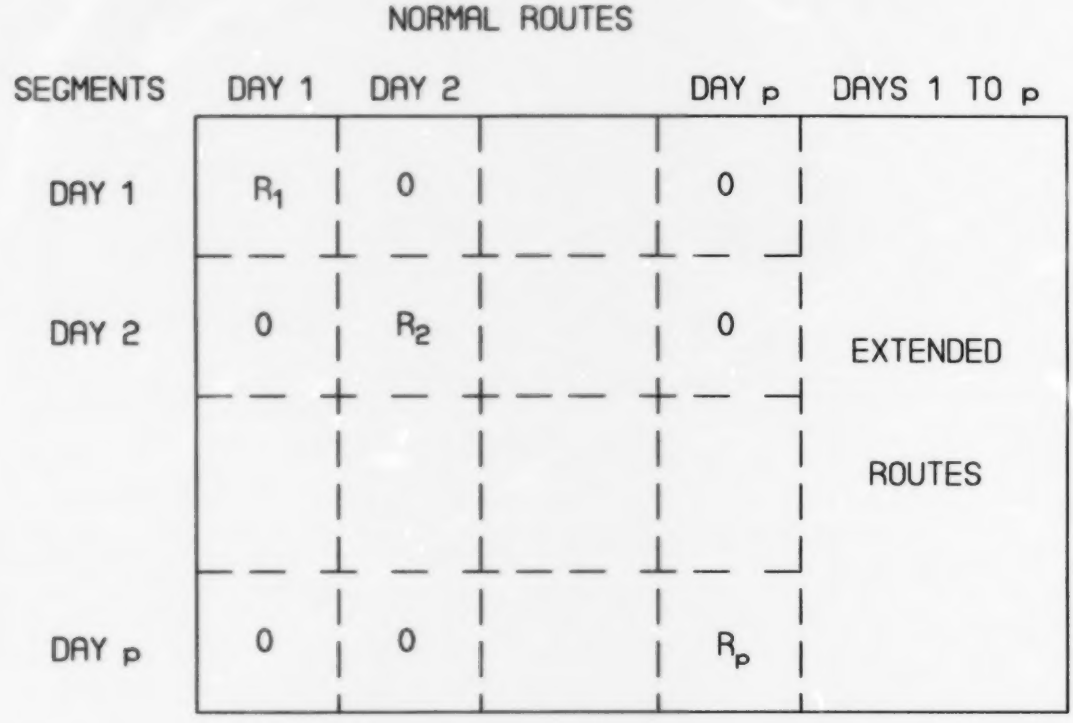

Figure 7. Structure of the Constraint Matrix

Figure 7 illustrates the fact that the set of all segments may be divided into p mutually exclusive subsets, where p represents the number of periods or "days" that normally makes up the planning period (usually 5):

$S_j = \{\ s \in S \mid$ segment starts at period $j,\ 1 \le j \le p\ \}$

$S = S_1 \cup S_2 \cup .... \cup S_p$ and $S_j \cap S_i = \emptyset\ \ \forall\ i,j\ \ 1 \le i,j \le p.$

The set of regular driver routes, R, may similarly be split into $p+1$ subsets R_1, R_2, ..., R_p, R_e, where

$$R_j = \{\ r \in R \mid \forall\, s \in \kappa_r,\ s \in S_j;\ \text{i.e., normal route}\ \}\quad 1 \le j \le p$$

$$R_e = \{\ r \in R \mid r \text{ extended route}\ \}.$$

Note $n_S = |S|$, the cardinality of set S, $n_{S_j} = |S_j|$, $1 \le j \le p$, $n_R = |R|$, and $n_{R_j} = |R_j|$ for $1 \le j \le p+1$.

Thus, the constraint matrix naturally decomposes into two components, one of which constitutes a block diagonal matrix, each block corresponding to the constraint matrix of a one-day-segments-normal-routes-only set covering problem. Therefore, since in the case of LTL freight transportation, normal routes may all be constructed quite easily (normal routes are significantly inferior in number to extended routes), and since, in general, normal routes are preferred to extended routes, we use the following algorithm to find the optimal solution to the set covering problem:

Step 1

For $j = 1, ..., p$

Solve the linear relaxation (PRHL) of the set-covering subproblem which corresponds to period j; use the constraint matrix A_j formed by the segments in S_j and the normal routes in R_j;

Let R_j^* be the set of routes which form the optimal solution to this problem and let Π_j^* be the corresponding simplex multipliers;

$R^* = \bigcup_j R_j^*$ and $\Pi^* = (\ \Pi_1^*\ \ \Pi_2^*\ \ \cdots\ \ \Pi_p^*\ )^T$.

Step 2

STOP if R^* covers all segments in S;

Otherwise

Note $\overline{S}$ as the set of segments not covered by routes in R^*.

Generate, by using the Π^* multipliers, extended routes, $\overline{R}$, to cover segments in $\overline{S}$.

Reoptimize; let R^* represent the set of routes in the optimal solution to the linear relaxation (PRHL).

Step 3

Extract, from R^*, R^I, the integer optimal solution to problem (PRH).

Note that, given the set of optimal simplex multipliers Π^*, the marginal cost of any route can be computed by using the "classical" linear programming formula

$$\overline{c}_r = c_r - \Pi^{*T} a_r$$

which, in the present context, is equivalent to

$$\overline{c}_r = c_r - \sum_{s \in \kappa_r} \pi_s^*$$

where π_s^* is the simplex multiplier associated to segment (constraint) s, $s \in S$.

It is therefore extremely easy to compute the reduced cost of a route, given its composition and the current optimal solution of the relaxed problem. Consequently, the relaxed subproblem may be solved very efficiently by a suitable combination of a good heuristic initial solution and the resolution of a (short: 2 or 3 repetitions only) series of linear programs, to which a limited number of good routes (i.e. with negative marginal costs) are admitted. Crainic and Roy (1990) discuss this procedure in detail. The same combination of post-optimization and column generation techniques is also used to generate extended routes, when required.

Finally, several alternatives exist to the problem of obtaining the integer optimal solution to the set covering problem, given the optimal solution to its linear relaxation. Exact algorithms, branch-and-bound, for example, may be used when exact solutions are required. When good approximate solutions are acceptable, we use an heuristic procedure (Salkin and Mathur, 1989; see also Crainic and Rousseau, 1987) which normally produces solutions very close (less than 1%) to the optimal ones.

VIII. EXPERIMENTAL RESULTS

Our experimentation was conducted on data from a large Canadian LTL motor carrier. The test network covered most of the provinces of Québec and Ontario and some parts of the Maritime provinces. The company data corresponds to the set of regular driver routes that the company operated in 1988, plus information concerning the timing of actual tractor movements - day and time of departure of each route occurrence. The company works on a five-day planning period and operates a network of twenty terminals, each located in a city which is also a driver base.

The available data thus already reflected the company's actual choices, policies and operational logic. Therefore, since data concerning the initial transportation plan was not available at the time, the experimentation fulfills more of a demonstration role for the worthiness of the method, rather than proving the full extent of its capabilities. Yet, even in this restrictive environment, the results are impressive.

Tables 3 and 4 summarize information concerning the routes operated by the company. Table 3 relates to routes. It displays the relative distribution of the 37 metaroutes that were used by the company and their frequency in the operated bids. It also shows the relative distribution of the 378 driver routes. Table 4 presents information concerning the distribution of segments according to the number of departures during the planning period. Data is displayed for the metasegments and the corresponding segments.

Type	Metaroutes		Driver Routes
	Relative Frequency	Frequency in Bids	Relative Frequency
Normal Extended	84% 16%	82.8% 17.2%	93.6% 6.4%
Total Number	37	100%	378

Table 3. Initial Route Distribution

It is worth noting that these figures justify our assertions concerning the limited usage of extended routes: more than 93% of all operated routes are made up of one duty period only. Another interesting feature is the fact that some 90% of all metasegments, which represent

over 93% of all segments, have at least five departures during a planning period. In fact, over two thirds of them (80% of the segments) have two daily departures, while one third (70% of the segments) have at least three departures daily. We also observed that, in general, the various departures of a metasegment during the planning period tend to be evenly distributed amongst the days, while departures of metasegments which "force" the use of extended routes are planned every other day. Finally, a majority of the operated bids (80% of them) prescribe that the driver takes the same route at each one of its departures. These figures amply illustrate the regularity of the operations which correspond to the stable part of demand, and empirically justify our approach.

Departures per Period	Number of Metasegments	Relative Frequency	Number of Segments	Relative Frequency
0 – 4	6	11%	24	2.71%
5 – 9	12	22%	66	7.45%
10 – 14	18	33%	184	20.77%
15 – 19	2	4%	30	3.39%
20 – 24	4	7%	88	9.93%
25 – 29	2	4%	50	5.64%
30 – 34	0	0%	0	0%
35 – 39	2	4%	72	8.13%
40 – 44	2	4%	80	9.03%
45 – 49	2	4%	90	10.16%
50 +	4	7%	202	22.80%
Total	54	100%	886	100%

Table 4. Segment Distribution in Initial Route Set

Two sets of segments have been generated for this data. The first has been directly extracted from the operated bids. Its segments' departure times are, therefore, those indicated in the company's documents. For the second set of segments, we generated departure times, following the procedure presented in Section IV, with special care toward insuring that the latest departure is such that the driver still arrives on time at destination.

Regular driver routes have been generated, and the optimal set of routes has been selected, by starting with each of these two segment sets. Crainic and Roy (1990) present a detailed analysis of the computational performance of the route generation and selection

procedures. The procedures were shown to be highly efficient and to produce a limited number of routes with "good" characteristics, in terms of route composition (number and arrangement of segments), number of tractor-without-trailer (TWT) movements, number of segments covered by more than one route, etc. In particular, Table 5 displays information about the routes generated and selected when using the second set of segments.

Stage	2-segment Routes	4-segment Routes	TWT Routes
Generation	4%	95%	12.5%
Selection	63%	35%	1.2%

Table 5. Characteristics of the Route Set

For the same run, 92.72% of the segments are covered by one route only, while 6.32% appear in two routes, and 0.95% are covered three times. These features are quite impressive for a set covering formulation. Besides, most of the cases where segments are covered more than once are due to imprecisions in the determination of the segment starting time and may easily be corrected. The other occurrences correspond to movements required to reposition tractors due to significant differences in departure frequencies in the opposite directions between two terminals. This phenomenon is similar to the deadheading practices observed in airline crew scheduling problems (Gershkoff, 1989; Crainic and Rousseau, 1987), and may similarly be addressed, by using post-optimal detection and adjustment procedures. Thus, a simple visual inspection of the solution permitted us to decrease by 41% the number of segments covered more than once. This phenomenon also underlines the interest for the interactive capabilities intended for the planning system.

It is also quite interesting to note that the solution of the relaxed problem offers very tight bounds for the integer formulation. Indeed, in some cases, the integer solution is directly obtained by solving the linear relaxation. Note that for our application, no post-optimization procedure is required to generate extended routes: all segments not covered by normal routes (5.4% of segments) formed "natural" pairs. Consequently, we manually formed the necessary extended routes.

Extremely interesting results may be derived by comparing the solutions obtained by our procedure to those of the initial data set.

The design of the routes is, in general, quite similar from one set to another, and several routes are common to both sets. This is quite natural, if one takes into account that the segments we use are rather directly derived from the routes operated by the company. Yet, one also notes that the application of our methodology has implied the selection of routes which are more efficient and more productive, and which make better use of the time available during a duty period. Thus, for example, two short routes, which are generated in all sets, have been combined in our optimal solution to result in a more efficient route.

This search for efficiency translates into significant improvements in route productivity measures for roughly the same total driver wages. In particular, the productivity cost decreased by some 17% (for the routes generated by using the second set of segments - for the first set, the figure was close to 15%), while the number of routes decreased by 13%.

These improvements resulted in 10% fewer drivers, and therefore trucks, required to satisfy demand. Thus, the gains obtained in designing regular driver routes are directly reflected in the utilization of the equipment and in the driver and rolling stock requirements. We also believe that even more significant results may be obtained by integrating the regular driver route design to the overall planning of operations, in particular to the tactical plan.

CONCLUSION

We have presented an integrated approach for the problem of designing regular driver routes in the less-than-truckload motor-carrier industry. Results obtained thus far clearly indicate that the potential savings which may be obtained by using this approach are quite substantial. Our methodology can provide assistance to motor carriers in making better use of their human and material resources, and hence enhance the overall planning process within the company and improve their productivity and competitivity.

ACKNOWLEDGMENTS

We wish to thank Ms. Lucie Bibeau for her excellent assistance in programming, and Mr. Alain Blanchette for his consistent and successful efforts in providing us with a very effective communication link with the industry. We also wish to express our gratitude to the following funding organizations: Le Fonds FCAR Québec, C.R.S.N.G. Canada, le ministère du Transport du Québec, and the Foundation of the Université du Québec à Montréal.

BIBLIOGRAPHY

Blanchette A. (1989), *Etablissement des routes et horaires réguliers dans les entreprises de transport routier des marchandises*, M.B.A. Thesis, département des Sciences administratives, Université du Québec à Montréal, publication #663, Centre de recherche sur les transports, Université de Montréal.

Crainic T.G. (1982), *Un modèle de planification tactique pour le transport ferroviaire des marchandises*, Ph.D. Thesis, publication #247, Centre de recherche sur les transports, Université de Montréal.

Crainic T.G. (1988), *Rail Tactical Planning: Issues, Models and Tools*, in *Freight Transport Planning and Logistics*, L. Bianco and A. La Bella (Eds.), 463-509, Springer-Verlag.

Crainic T.G. and Rousseau J.-M. (1986), *Multicommodity, Multimode Freight Transportation: A General Modeling and Algorithmic Framework for the Service Network Design Problem*, **Transportation Research B** 20(3), 225-242.

Crainic T.G. and Rousseau J.-M. (1987), *The Column Generation Principle and the Airline Crew Scheduling Problem*, **INFOR** 25 (2), 136-151.

Crainic T.G. and Roy J. (1988), *O.R. Tools for Tactical Freight Transportation Planning*, **European Journal of Operations Research** 33 (3), 290-297.

Crainic T.G. and Roy J. (1990), *Une approche de recouvrement d'ensembles pour l'établissement d'horaires de chauffeurs dans le transport routier de charges partielles*, **R.A.I.R.O. – Recherche Opérationnelle** 2, 1-36.

Crainic T.G., Roy J. and Delorme L. (1989), *Le problème de l'établissement des routes–horaires dans le transport routier de charges partielles*, publication #636, Centre de recherche sur les transports, Université de Montréal.

Delorme L. Roy J. and Rousseau J.M. (1988), *Motor–Carriers Operations Planning Models: A State of the Art*, in *Freight Transport Planning and Logistics* L. Bianco and A. La Bella (Eds.), 510-545, Springer-Verlag.

Desrochers M. (1986), *La fabrication d'horaires de travail pour les conducteurs d'autobus par une méthode de génération de colonnes*, Ph.D. Thesis, publication #470, Centre de recherche sur les transports, Université de Montréal.

Gershkoff I. (1989), *Optimizing Flight Crew Schedules*, **Interfaces** 19(4), 29-43.

Lavoie S., Minoux M. and Odier E. (1988), *A New Approach for Crew Pairing Problems by Column Generation with an Application to Air Transportation*, **European Journal of Operational Research** 35, 45-58.

Rousseau J.-M. (1984), *Crew Scheduling Methods in the Operation and Planning of Transportation Systems*, in *Transportation Planning Models*, M. Florian (Ed.), 439-474, North-Holland.

Roy J. (1984), *Un modèle de planification globale pour le transport routier des marchandises*, Ph.D. Thesis, publication #402, Centre de recherche sur les transports, Université de

Montréal.

Roy J. and Delorme L. (1989), *NETPLAN: A Network Optimization Model for Tactical Planning in the Less-than-Truckload Motor-Carrier Industry*, **INFOR** 27 (1), 22-35.

Falkin H.M. and Mathur K. (1989), *Fundamentals of Integer Programming*, North-Holland, New York, U.S.A.